

DETAILED ACTION

1. This Action is in response to Application Number 10/683,933 received on 7/26/2008.
2. Claims 1-8, 14, 19-20, 23-28, 33, 36-39, 41-51, 53-54, 56-61 are presented for examination.

Response to Arguments

Applicant's arguments filed 07/26/2008 have been fully considered but they are not persuasive.

Objection to the Specification

Applicant is respectfully requested to correct the "Cross Reference to Related Applications" section of Applicant's Specification for being incomplete as indicated in the previous office action. Correction is required. See MPEP § 608.01(b).

Rejections under 35 U.S.C. 102

DiMambro Reference

Applicant asserts, "DiMambro falls under 102(e)/103(c) exception. Accordingly, the Office is requested to remove DiMambro as 102(e) art against the pending claims [Response, page 9].

Examiner respectfully disagrees.

The DiMambro reference qualifies as anticipatory prior art under 35 U.S.C 102, and therefore does not fall within the 35 U.S.C. 103(c) exception.

Applicant is requested to review MPEP 706.02(l)(1) which states, “**35 U.S.C. 103(c) applies only to prior art usable in an obviousness rejection under 35 U.S.C. 102**”

103. Subject matter that qualifies as anticipatory prior art under 35 U.S.C. 102 is not affected, and may still be used to reject claims as being anticipated.”

The DiMambro reference was not used in a 103(a) rejection and is therefore not affected by the 103(c) exception. As such, the rejection is respectfully maintained.

Chang Reference

Applicant asserts, “As can be clearly seen, in Chang the data packets for a particular connection are distributed among the various threads running in various CPUs available on the system. This is different from the claimed invention wherein the data packets for a particular connection are all queued to a single queue associated with a single processor” [Response, page 9].

Examiner respectfully disagrees.

As explicitly shown in Figure 3, the setup of Chang shows that each queue (62, 64, 66, 68) has its own respective CPU (54, 56, 58, 60). As also explicitly shown in Figure 3, packets enter the queue through a HASH 50. As explicitly stated in Chang, “Packets are distributed to one of the N queues by using a hashing function based on the source MAC address, source IP address, or the packet’s source and destination TCP port number, or all or a combination of the foregoing” (Chang, col. 2, lines 64-67). Chang further explicitly states, “It is a feature of the invention to hash, in accordance with the hash function 50 shown in Fig. 3, each of these MACS into the aforementioned

plurality of queues 62-68, whereby a hash is provided by the hash function 50 to packets in the device drivers 42 by means of the interconnection 52 will occur for given devices to the network so that packets associated with a given device will be handled in the same particular queue.

It is well known to one of ordinary skill in the art that packets of the same connection will have the same addresses. Since the teachings of Chang distribute the packets to a particular queue based on a hashing function of these addresses, then all packets having the same addresses will be sent to the same queue. Therefore, all packets of the same connection will be sent to the same queue.

As such, the rejection is respectfully maintained.

Rejections under U.S.C. 103(a)

Applicant asserts, “it would not be obvious to one of ordinary skill in the art to substitute the queue of Chang with any type of queue, including a squeue, to arrive at the claimed invention” [Response, page 10].

Examiner respectfully disagrees.

The Chang reference clearly disclosed the functionality performed by the queue as recited in claim 3. The limitation, “wherein said processing queue is an squeue”, does nothing to change the actual functionality of the queue. The squeue still performs the same functionality as described in claim 3. It is clearly obvious to one of ordinary skill in the art to substitute one device for another device in order to achieve the same result. In this case, it would have been obvious to one of ordinary skill to substitute the

queue of Chang with any type of queue that can perform the same functionality as the queue of Chang. Such a conclusion for obviousness follows the exemplary rationales provided in MPEP 2143. As such, the rejection is respectfully maintained.

Applicant reiterates the argument as presented in the 102 section of Applicant's Response [Response, page 10]

Examiner respectfully disagrees.

For the reasons set forth above, the rejection is respectfully maintained.

Applicant's final argument includes, "Combining Syvanne with Chang will not teach all the limitations of the claimed invention, which is to assign all packets from a connection to a single queue associated with a single thread of a single processor within a multi-process server system" [Response, page 11].

Examiner respectfully disagrees for the same reasons set forth indicated in the response to Applicant's arguments with respect to the 102 rejections, detailed above. It is evident from the above explanation and the rejection that Chang disclosed these features of the claims. As such, the rejection is respectfully maintained.

It is the Examiner's position that Applicant has not yet submitted claims drawn to limitations, which define the operation and apparatus of Applicant's disclosed invention in manner, which distinguishes over the prior art.

Failure for Applicant to significantly narrow definition/scope of the claims and supply arguments commensurate in scope with the claims implies the Applicant intends broad interpretation be given to the claims. The Examiner has interpreted the claims with scope parallel to the Applicant in the response and reiterates the need for the Applicant to more clearly and distinctly define the claimed invention.

Specification

3. The Specification is objected to because the “Cross Reference to Related Applications” section is incomplete. Correction is required. See MPEP § 608.01(b).

Claim Rejections - 35 USC § 102

The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(b) the invention was patented or described in a printed publication in this or a foreign country or in public use or on sale in this country, more than one year prior to the date of application for patent in the United States.

(e) the invention was described in a patent granted on an application for patent by another filed in the United States before the invention thereof by the applicant for patent, or on an international application by another who has fulfilled the requirements of paragraphs (1), (2), and (4) of section 371(c) of this title before the invention thereof by the applicant for patent.

The changes made to 35 U.S.C. 102(e) by the American Inventors Protection Act of 1999 (AIPA) and the Intellectual Property and High Technology Technical Amendments Act of 2002 do not apply when the reference is a U.S. patent resulting directly or indirectly from an international application filed before November 29, 2000.

Therefore, the prior art date of the reference is determined under 35 U.S.C. 102(e) prior to the amendment by the AIPA (pre-AIPA 35 U.S.C. 102(e)).

4. Claims 1-4, 20, 27-28, 33, 36-37, 39, 41-47, 53-54, and 60-61 are rejected under 35 U.S.C. 102(e) as being anticipated by DiMambro (U.S. 7076545).

5. Regarding claim 1, DiMambro disclosed a method for processing packets through a plurality of protocol layers comprising:

accessing a packet associated with a connection (DiMambro, col. 1, lines 39-40); and

processing said packet through said plurality of protocol layers using a single thread from a single processor by assigning said connection to a single processor of a multiprocessor server system for processing wherein packets associated with said connection are directed to said single thread in said single processor for processing and wherein connection state information used by said plurality of protocol layers is preserved by mutual exclusion of other threads processing packets for said connection through said plurality of protocol layers (DiMambro, col. 1, lines 41-50. col. 3, lines 3-15, each flow may be assigned to a service queue that corresponds to a processor for protocol processing, col. 5, lines 5-7, management of each service queue is independently performed of the other queues).

6. Regarding claim 2, DiMambro disclosed the limitations as described in claim 1, including wherein said single thread is uninterrupted while processing said packet through said plurality of protocol layers (DiMambro, col. 2, lines 49-58).

7. Regarding claim 3, DiMambro disclosed the limitations as described in claim 1, including assigning said packet to a processing queue wherein said processing queue provides single threaded processing of said packet through said plurality of protocol layers (DiMambro, col. 1, lines 41-49).

8. Regarding claim 4, DiMambro disclosed the limitations as described in claim 3, including wherein said processing queue provides single threaded processing of said packet through said plurality of protocol layers by assigning only one packet to be processed by said plurality of protocol layers at a time (DiMambro, col. 2, lines 49-58).

9. Regarding claim 20, DiMambro disclosed a method for processing packets comprising:

accessing a packet associated with a connection (DiMambro, col. 1, lines 39-40); and

assigning said packet to a processing queue associated with a single processor of a multiprocessor server system wherein said processing queue provides uninterrupted single threaded processing of said data packet associated with the connection through a plurality of protocol layers using a single thread of the single

processor, wherein state information of the packet within the connection is preserved so as to mutually exclude other threads from processing packets of said connection through said plurality of protocol layers (DiMambro, col. 1, lines 35-50. col. 3, lines 3-15, each flow may be assigned to a service queue that has a thread and corresponds to a processor for protocol processing, col. 5, lines 5-7, management of each service queue is independently performed of the other queues).

10. Regarding claim 27, DiMambro disclosed the limitations as described in claim 20, including wherein said plurality of protocol layers includes a TCP protocol layer (DiMambro, col. 2, lines 49-55).

11. Regarding claim 28, DiMambro disclosed the limitations as described in claim 20, including wherein said plurality of protocol layers includes an IP protocol layer (DiMambro, col. 2, lines 49-55).

12. Regarding claim 33, DiMambro disclosed the limitations as described in claim 23, including wherein said queue is associated with said connection data structure (DiMambro, col. 5, lines 18-25).

13. Regarding claim 36, DiMambro disclosed a multiprocessor server system comprising:

a plurality of processors for processing packets through a plurality of protocol layers (DiMambro, col. 1, lines 35-40);

a plurality of threads running in the plurality of processors (DiMambro, col. 1, lines 35-45);

a plurality of queues, each queue associated with a respective processor of said plurality of processors (DiMambro, col. 1, lines 41-45); and

a memory resident connection data structure for assigning packets of a connection to a queue of said plurality of queues for processing said packets of said connection using a single thread associated with the single queue of a corresponding processor of said plurality of processors (DiMambro, col. 1, lines 32-45, col. 4, lines 8-20, col. 5, lines 15-25).

14. Regarding claim 37, DiMambro disclosed the limitations as described in claim 36, including wherein said connections are TCP connections (DiMambro, col. 4, lines 49-51).

15. Regarding claim 39, DiMambro disclosed the limitations as described in claim 36, including wherein a processor of said plurality of processors processes a packet of its queue without interruption through said plurality of protocol layers except for scheduling another packet on its queue (DiMambro, col. 5, lines 1-7).

16. Regarding claim 41, DiMambro disclosed the limitations as described in claim 37, including wherein said connection data structure is established for a new connection upon receiving a new connection request and wherein said connection data structure comprises an identifier of a queue associated with the single thread of the same corresponding processor to which all packets of said new connection are to be assigned (DiMambro, col. 1, lines 32-45, col. 4, lines 8-20, col. 5, lines 15-25).

17. Regarding claim 42, DiMambro disclosed the limitations as described in claim 36, including a plurality of cache memories, each cache associated with a respective processor of said plurality of processors (DiMambro, col. 1, lines 39-45).

18. Regarding claim 43, DiMambro disclosed the limitations as described in claim 36, including wherein state information of any given packet of a same connection is preserved because said packets of said same connection are individually mutually excluded from said protocol layers (DiMambro, col. 1, lines 45-50).

19. Regarding claim 44, DiMambro disclosed a computer system comprising a processor coupled to a bus and a memory coupled to said bus and comprising instructions (DiMambro, col. 2, lines 24-35) that when executed implement a method for processing data packets comprising:
accessing a packet associated with a connection (DiMambro, col. 1, lines 39-40);
and

processing said packet through said plurality of protocol layers using a single thread from a single processor by assigning said connection for processing to a single processor of a multiprocessor server system wherein packets associated with said connection are directed to said single thread associated with said single processor for processing, wherein connection state information used by said plurality of protocol layers is preserved by mutual exclusion of other threads processing packets for said connection through said plurality of protocol layers (DiMambro, col. 1, lines 35-50. col. 3, lines 3-15, each flow may be assigned to a service queue that has a thread and corresponds to a processor for protocol processing, col. 5, lines 5-7, management of each service queue is independently performed of the other queues).

20. Regarding claim 45, DiMambro disclosed the limitations as described in claim 44, including wherein said single thread is uninterrupted while processing said packet through said plurality of protocol layers (DiMambro, col. 2, lines 43-47).

21. Regarding claim 46, DiMambro disclosed the limitations as described in claim 44, including wherein said packet are assigned to a processing queue wherein said processing queue provides single threaded processing of said packet through said plurality of protocol layers (DiMambro, col. 1, lines 38-45, col. 2, lines 43-47).

22. Regarding claim 47, DiMambro disclosed the limitations as described in claim 46, including wherein said processing queue provides single threaded processing of said

packet through said plurality of protocol layers by assigning only one packet to be processed by said plurality of protocol layers at a time (DiMambro, col. 2, lines 49-58).

23. Regarding claim 53, DiMambro disclosed a computer system comprising a processor coupled to a bus and a memory coupled to said bus and comprising instructions (DiMambro, col. 2, lines 24-35) that when executed implement a method for processing data packets comprising:

accessing a packet associated with a connection (DiMambro, col. 1, lines 39-40); and

assigning said packet to a processing queue associated with a single processor of a multi processor server system, wherein said processing queue provides uninterrupted single threaded processing of said data packet associated with the connection through a plurality of protocol layers using a single thread of the single processor, wherein state information of the packet within the connection is preserved so as to mutually exclude other threads from processing packets of said connection through said plurality of protocol layers (DiMambro, col. 1, lines 35-50. col. 3, lines 3-15, each flow may be assigned to a service queue that has a thread and corresponds to a processor for protocol processing, col. 5, lines 5-7, management of each service queue is independently performed of the other queues).

24. Regarding claim 54, DiMambro disclosed the limitations as described in claim 53, including wherein said processing queue provides mutual exclusion of same-connection packet processing through said plurality of protocol layers (DiMambro, col. 5, lines 5-7).

25. Regarding claim 60, DiMambro disclosed the limitations as described in claim 53, including wherein said plurality of protocol layers includes a TCP protocol layer (DiMambro, col. 4, lines 49-51).

26. Regarding claim 61, DiMambro disclosed the limitations as described in claim 53, including wherein said plurality of protocol layers includes an IP protocol layer (DiMambro, col. 4, lines 49-51).

27. Claims 1-5, 20, 27-28, 44-48, 53-54, 60-61 are rejected under 35 U.S.C. 102(b) as being anticipated by Chang et al. (U.S. 6,338,078).

28. Regarding claim 1, Chang disclosed a method for processing packets through a plurality of protocol layers comprising:

accessing a packet associated with a connection (Chang, col. 2, lines 56-58);
and

processing said packet through said plurality of protocol layers using a single thread from a single processor by assigning said connection to a single processor of a multiprocessor server system for processing wherein packets associated with said

connection are directed to said single thread in said single processor for processing and wherein connection state information used by said plurality of protocol layers is preserved by mutual exclusion of other threads processing packets for said connection through said plurality of protocol layers (Chang, Fig. 3, queues 62,64,66, 68 having corresponding processors CPU's 54,56,58,60; col. 2, lines 55-60, the packets are distributed to N high priority threads; col. 2, lines 64-67, Packets are distributed by a hashing function, hashing addresses of the packets to determine which queue/CPU pair it should be sent to, ; col. 4, lines 40-45, protocol processing; col. 4, line 60, single thread processing col. 5, lines 60-65; col. 6, lines 45-50, each CPU handles a particular queue).

29. Regarding claim 2, Chang disclosed the limitations as described in claim 1, including wherein said single thread is uninterrupted while processing said packet through said plurality of protocol layers (Chang, col. 5, lines 30-33, lines 60-66).

30. Regarding claim 3, Chang disclosed the limitations as described in claim 1, including assigning said packet to a processing queue wherein said processing queue provides single threaded processing of said packet through said plurality of protocol layers (Chang, col. 6, lines 44-47).

31. Regarding claim 4, Chang disclosed the limitations as described in claim 3, including wherein said processing queue provides single threaded processing of said

packet through said plurality of protocol layers by assigning only one packet to be processed by said plurality of protocol layers at a time (Chang, col. 6, lines 45-50).

32. Regarding claim 5, Chang disclosed the limitations as described in claim 4, including wherein said packet is assigned to said processing queue based on address information of said connection (Chang, col. 2, lines 64-67).

33. Regarding claim 20, Chang disclosed a method for processing packets comprising:

accessing a packet associated with a connection (Chang, col. 2, lines 56-58); and

assigning said packet to a processing queue associated with a single processor of a multiprocessor server system wherein said processing queue provides uninterrupted single threaded processing of said data packet associated with the connection through a plurality of protocol layers using a single thread of the single processor, wherein state information of the packet within the connection is preserved so as to mutually exclude other threads from processing packets of said connection through said plurality of protocol layers (Chang, Fig. 3, queues 62,64,66, 68 having corresponding processors CPU's 54,56,58,60; col. 2, lines 55-60, the packets are distributed to N high priority threads; col. 2, lines 60-67, Packets are distributed by a hashing function, hashing addresses of the packets to determine which queue/CPU pair it should be sent to, ; col. 4, lines 40-45, protocol processing; col. 4, line 60, single

thread processing col. 5, lines 60-65; col. 6, lines 45-50, each CPU handles a particular queue).

34. Regarding claim 27, Chang disclosed the limitations as described in claim 20, including wherein said plurality of protocol layers includes a TCP protocol layer (Chang, col. 4, lines 40-45).

35. Regarding claim 28, Chang disclosed the limitations as described in claim 20, including wherein said plurality of protocol layers includes an IP protocol layer (Chang, Fig. 4, 74).

36. Regarding claim 44, Chang disclosed a computer system comprising a processor coupled to a bus and a memory coupled to said bus and comprising instructions (Chang, col. 3, lines 25-55) that when executed implement a method for processing data packets comprising:

accessing a packet associated with a connection (Chang, col. 2, lines 56-58); and

processing said packet through said plurality of protocol layers using a single thread from a single processor by assigning said connection for processing to a single processor of a multiprocessor server system wherein packets associated with said connection are directed to said single thread associated with said single processor for processing, wherein connection state information used by said plurality of protocol

layers is preserved by mutual exclusion of other threads processing packets for said connection through said plurality of protocol layers (Chang, Fig. 3, queues 62,64,66, 68 having corresponding processors CPU's 54,56,58,60; col. 2, lines 55-60, the packets are distributed to N high priority threads; col. 2, lines 64-67, Packets are distributed by a hashing function, hashing addresses of the packets to determine which queue/CPU pair it should be sent to, ; col. 4, lines 40-45, protocol processing; col. 4, line 60, single thread processing col. 5, lines 60-65; col. 6, lines 45-50, each CPU handles a particular queue).

37. Regarding claim 45, Chang disclosed the limitations as described in claim 44, including wherein said single thread is uninterrupted while processing said packet through said plurality of protocol layers (Chang, col. 5, lines 30-33, lines 60-66)).

38. Regarding claim 46, Chang disclosed the limitations as described in claim 44, including wherein said packet are assigned to a processing queue wherein said processing queue provides single threaded processing of said packet through said plurality of protocol layers (Chang, col. 4, lines 60-67).

39. Regarding claim 47, Chang disclosed the limitations as described in claim 46, including wherein said processing queue provides single threaded processing of said packet through said plurality of protocol layers by assigning only one packet to be processed by said plurality of protocol layers at a time (Chang, col. 4, lines 60-67).

40. Regarding claim 48, hang disclosed the limitations as described in claim 47, including wherein said packet is assigned to said processing queue based on address information of said connection (Chang, col. 2, lines 64-67).

41. Regarding claim 53, Chang disclosed a computer system comprising a processor coupled to a bus and a memory coupled to said bus and comprising instructions (Chang, col.) that when executed implement a method for processing data packets comprising:

accessing a packet associated with a connection (Chang, col. 2, lines 56-58); and

assigning said packet to a processing queue associated with a single processor of a multi processor server system, wherein said processing queue provides uninterrupted single threaded processing of said data packet associated with the connection through a plurality of protocol layers using a single thread of the single processor, wherein state information of the packet within the connection is preserved so as to mutually exclude other threads from processing packets of said connection through said plurality of protocol layers (Chang, Fig. 3, queues 62,64,66, 68 having corresponding processors CPU's 54,56,58,60; col. 2, lines 55-60, the packets are distributed to N high priority threads; col. 2, lines 64-67, Packets are distributed by a hashing function, hashing addresses of the packets to determine which queue/CPU pair it should be sent to, ; col. 4, lines 40-45, protocol processing; col. 4, line 60, single

thread processing col. 5, lines 60-65; col. 6, lines 45-50, each CPU handles a particular queue).

42. Regarding claim 54, Chang disclosed the limitations as described in claim 53, including wherein said processing queue provides mutual exclusion of same-connection packet processing through said plurality of protocol layers (Chang, col. 4, lines 60-67).

43. Regarding claim 60, Chang disclosed the limitations as described in claim 53, including wherein said plurality of protocol layers includes a TCP protocol layer (Chang, col. 4, lines 40-45).

44. Regarding claim 61, Chang disclosed the limitations as described in claim 53, including wherein said plurality of protocol layers includes an IP protocol layer (Chang, col. 4, lines 40-45).

Claim Rejections - 35 USC § 103

The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the manner in which the invention was made.

45. Claim 14 is rejected under 35 U.S.C. 103(a) as being unpatentable over Chang et al. (U.S. 6,338,078).

46. Regarding claim 14, Chang disclosed the limitations as described in claim 3. Chang did not explicitly state wherein the queue is an squeue.

However, it would have been obvious to one of ordinary skill in the art at the time the invention was made to substitute the queue of Chang with any type of queue, including an squeue, as long as the same functionality is performed, i.e. in order to achieve the predictable result of the queue queuing packets for processing. The type of queue is a matter of design choice since the same functionality is performed.

47. Claims 6-8, 19, 23-26, 33, 36-39, 41-43, 49-51, and 56-59 are rejected under 35 U.S.C. 103(a) as being unpatentable over Chang et al. (U.S. 6,338,078) in view of Syvanne (U.S. 2002/0112188).

48. Regarding claims 6, 23, 49, and 56, Chang disclosed the limitations as described in claims 1, 20, 44, and 53. Chang also disclosed that packets are distributed to one of the N queues by using a hashing function (Chang, col. 2, lines 64-66).

Chang did not explicitly state generating a unique connection data structure specific to said connection stored in said packet associated with the connection.

In an analogous art, Syvanne disclosed a method for handling information about packet data connections in which connection data structures are created (Syvanne,

[0009]). As such, Syvanne suggests the use of data structures for classifying packets and shows that such packet classification was well known at the time the invention was made.

Therefore, it would have been obvious to one of ordinary skill in the art at the time the invention was made to incorporate using data structures to classify packets according to connections into the system of Chang in order to obtain the predictable result of distributing packets to its corresponding queue/processor duo based on the connection that the packet belongs, thereby providing a more scalable system allowing multiple ways of classifying packets.

49. Regarding claims 7, 24, 50, and 57, Chang and Syvanne disclosed the limitations as described in claims 6, 23, 49, and 56, including wherein said address information comprises a local IP address and a remote IP address (Chang, Fig. 4, 74). See motivation above.

50. Regarding claims 8, 25, 51, 58, Chang and Syvanne disclosed the limitations as described in claims 7, 24, 49, 57, including wherein said address information further comprises a remote port address and a local port address (Chang, Fig. 4, 76). See motivation above.

51. Regarding claims 19, 26, 59, Chang and Syvanne disclosed the limitations as described in claim 6, 25, and 58. Chang and Syvanne did not explicitly state wherein

subsequent data packets of said connection are assigned to said single processor based on said connection data structure.

However as shown above, Syvanne suggests the use of data structures for classifying packets and shows that such packet classification was well known at the time the invention was made.

Therefore, it would have been obvious to one of ordinary skill in the art at the time the invention was made to incorporate using data structures to classify packets according to connections into the system of Chang in order to obtain the predictable result of distributing packets to its corresponding queue/processor duo based on the connection that the packet belongs, thereby providing a more scalable system allowing multiple ways of classifying packets.

52. Regarding claim 33, Chang and Syvanne disclosed the limitations as described in claim 23. Chang and Syvanne did not explicitly state wherein said queue is associated with said connection data structure.

However, it would have been obvious to one of ordinary skill in the art at the time the invention was made that using the connection data structure to properly classify the packets and properly distribute them to the queue/processor duo includes an association between the connection data structure and the queue. See motivation above.

53. Regarding claim 36, Chang disclosed a multiprocessor server system comprising:

a plurality of processors for processing packets through a plurality of protocol layers (Chang, col. 2, lines 55-56);

a plurality of threads running in the plurality of processors (Chang, col. 2, lines 55-60, “N high priority threads”); and

a plurality of queues, each queue associated with a respective processor of said plurality of processors (Chang, col. 2, lines 59-65).

Chang also disclosed assigning packets of a connection to a queue of said plurality of queues for processing said packets of said connection using a single thread associated with the single queue of a corresponding processor of said plurality of processors (Chang, Fig. 3, queues 62,64,66, 68 having corresponding processors CPU’s 54,56,58,60; col. 2, lines 55-60, the packets are distributed to N high priority threads; col. 2, lines 64-67, Packets are distributed by a hashing function, hashing addresses of the packets to determine which queue/CPU pair it should be sent to, ; col. 4, lines 40-45, protocol processing; col. 4, line 60, single thread processing col. 5, lines 60-65; col. 6, lines 45-50, each CPU handles a particular queue).

Chang did not explicitly state including a memory resident connection data structure for such assignment.

In an analogous art, Syvanne disclosed a method for handling information about packet data connections in which connection data structures are created (Syvanne, [0009]). As such, Syvanne suggests the use of data structures for classifying packets

and shows that such packet classification was well known at the time the invention was made.

Therefore, it would have been obvious to one of ordinary skill in the art at the time the invention was made to incorporate using data structures to classify packets according to connections into the system of Chang in order to obtain the predictable result of distributing packets to its corresponding queue/processor duo based on the connection that the packet belongs, thereby providing a more scalable system allowing multiple ways of classifying packets.

54. Regarding claim 37, Chang and Syvanne disclosed the limitations as described in claim 36, including wherein the connections are TCP connections (Chang, col. 2, lines 65-67). See motivation above.

55. Regarding claim 38, Chang and Syvanne disclosed the limitations as described in claim 37, including wherein said plurality of protocol layers comprise IP, TCP, and socket layers (Chang, Fig. 4). See motivation above.

56. Regarding claim 39, Chang and Syvanne disclosed the limitations as described in claim 36, including wherein a processor of said plurality of processors processes a packet of its queue without interruption through said plurality of protocol layers except for scheduling another packet on its queue (Chang, col. 2, lines 55-67). See motivation above.

57. Regarding claim 41, Chang and Syvanne disclosed the limitations as described in claim 37, including wherein said connection data structure is established for a new connection upon receiving a new connection request and wherein said connection data structure comprises an identifier of a queue associated with the single thread of the same corresponding processor to which all packets of said new connection are to be assigned (See rejection for claim 19, see also Chang, Fig. 3, queues 62,64,66, 68 having corresponding processors CPU's 54,56,58,60; col. 2, lines 55-60, the packets are distributed to N high priority threads; col. 2, lines 60-67, Packets are distributed by a hashing function, hashing addresses of the packets to determine which queue/CPU pair it should be sent to, ; col. 4, lines 40-45, protocol processing; col. 4, line 60, single thread processing col. 5, lines 60-65; col. 6, lines 45-50, each CPU handles a particular queue). See motivation above.

58. Regarding claim 42, Chang and Syvanne disclosed the limitations as described in claim 36, including a plurality of cache memories, each cache associated with a respective processor of said plurality of processors (Chang, Fig. 3, 62, 64, 66, 68). See motivation above.

59. Regarding claim 43, Chang and Syvanne disclosed the limitations as described in claim 36, including wherein state information of any given packet of a same connection is preserved because said packets of said same connection are individually

mutually excluded from said protocol layers (Chang, col. 2, lines 55-67). See motivation above.

Conclusion

Examiner's Note: Examiner has cited particular columns and line numbers in the references applied to the claims above for the convenience of the applicant. Although the specified citations are representative of the teachings of the art and are applied to specific limitations within the individual claim, other passages and figures may apply as well. It is respectfully requested from the applicant in preparing responses, to fully consider the references in entirety as potentially teaching all or part of the claimed invention, as well as the context of the passage as taught by the prior art or disclosed by the Examiner.

In the case of amending the claimed invention, Applicant is respectfully requested to indicate the portion(s) of the specification which dictate(s) the structure relied on for proper interpretation and also to verify and ascertain the metes and bounds of the claimed invention.

THIS ACTION IS MADE FINAL. Applicant is reminded of the extension of time policy as set forth in 37 CFR 1.136(a).

A shortened statutory period for reply to this final action is set to expire THREE MONTHS from the mailing date of this action. In the event a first reply is filed within TWO MONTHS of the mailing date of this final action and the advisory action is not mailed until after the end of the THREE-MONTH shortened statutory period, then the shortened statutory period will expire on the date the advisory action is mailed, and any

extension fee pursuant to 37 CFR 1.136(a) will be calculated from the mailing date of the advisory action. In no event, however, will the statutory period for reply expire later than SIX MONTHS from the mailing date of this final action.

Any inquiry concerning this communication or earlier communications from the examiner should be directed to J. Bret Dennison whose telephone number is (571) 272-3910. The examiner can normally be reached on M-F 8:30am-5pm.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Tonia Dollinger can be reached on (571) 272-4170. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

/J. Bret Dennison/
Examiner, Art Unit 2443